


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before March 2004

Found 1,812 of 151,748

Terms used vertical horizontal database

Sort results by

☒ Save results to a Binder

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results

☒ Search Tips

☐ Open results in a new window

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

21 [Historical queries along multiple lines of time evolution](#)

Gad M. Landau, Jeanette P. Schmidt, Vassilis J. Tsotras

October 1995 **The VLDB Journal — The International Journal on Very Large Data****Bases**, Volume 4 Issue 4**Publisher:** Springer-Verlag New York, Inc.Full text available: [pdf\(1.41 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Traditional approaches to addressing historical queries assume a *single* line of time evolution; that is, a system (database, relation) evolves over time through a sequence of transactions. Each transaction always applies to the unique, current state of the system, resulting in a new current state. There are, however, complex applications where the system's state evolves into *multiple* lines of evolution. In general, this creates a tree (hierarchy) of evolution lines, where each tree ...

Keywords: CAD databases, access methods, data-structures, rollback databases

22 [Data file management in shift-register memories](#)

Werner E. Kluge

June 1978 **ACM Transactions on Database Systems (TODS)**, Volume 3 Issue 2**Publisher:** ACM PressFull text available: [pdf\(1.22 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The paper proposes a shift-register memory, structured as a two-dimensional array of uniform shift-register loops which are linked by flow-steering switches, whose switch control scheme is tailored to perform with great efficiency data management operations on sequentially organized files. The memory operates in a linear input/output mode to perform record insertion, deletion, and relocation on an existing file, and in a sublinear mode for rapid internal file movement to expedite file posit ...

Keywords: LIFO/FIFO operation modes, data transformations, deletion, insertion, management of sequentially organized files, record retrieval, relocation, shift-register memories, updating

23 [Integrating relational databases with support for updates](#)

M. Samy Gamal-Eldin, G. Thomas, R. Elmasri

January 2000 **Proceedings of the first international symposium on Databases in parallel and distributed systems DPDS '88**

Publisher: IEEE Computer Society Press

Full text available:  pdf(1.00 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Most work on database integration has considered only support for data retrieval, not support for updates, and often the use of a special semantically rich data model has been required. In this paper we present an approach to database integration which supports updates and which uses only the standard relational data model. Many of the ideas used in this approach are applicable to database integration in the context of other data models as well.


24 Expert design tools for physical database design



Rajiv Tewari

September 1990 **Proceedings of the 1990 ACM SIGBDP conference on Trends and directions in expert systems SIGBDP '90**

Publisher: ACM Press

Full text available:  pdf(926.35 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

25 Options in physical database design



Goetz Graefe

September 1993 **ACM SIGMOD Record**, Volume 22 Issue 3

Publisher: ACM Press

Full text available:  pdf(812.87 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

A cornerstone of modern database systems is physical data independence, i.e., the separation of a type and its associated operations from its physical representation in memory and on storage media. Users manipulate and query data at the logical level; the DBMS translates these logical operations to operations on files, indices, records, and disks. The efficiency of these physical operations depends very much on the choice of data representations. Choosing a physical representatio ...

26 Horizontal partitioning



Jan E. Bond

February 1988 **Proceedings of the 1988 ACM sixteenth annual conference on Computer science CSC '88**

Publisher: ACM Press

Full text available:  pdf(580.05 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)


27 View indexing in relational databases



Nicholas Roussopoulos

June 1982 **ACM Transactions on Database Systems (TODS)**, Volume 7 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.94 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The design and maintenance of a useful database system require efficient optimization of the logical access paths which demonstrate repetitive usage patterns. Views (classes of queries given by a query model) are an appropriate intermediate logical representation for database. Frequently accessed views of databases need to be supported by indexing to enhance retrieval. This paper investigates the problem of selecting an optimal index set of views and describes an efficient algorithm for thi ...

Keywords: index selection

28 Database partitioning in a cluster of processors



Domenico Sacca, Gio Wiederhold

March 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 1

Publisher: ACM Press

Full text available: pdf(2.39 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In a distributed database system the partitioning and allocation of the database over the processor nodes of the network can be a critical aspect of the database design effort. In this paper we develop and evaluate algorithms that perform this task in a computationally feasible manner. The network we consider is characterized by a relatively high communication bandwidth, considering the processing and input output capacities in its processors. Such a balance is typical if the processors are ...

29 A homogeneous relational model and query languages for temporal databases



Shashi K. Gadia

October 1988 **ACM Transactions on Database Systems (TODS)**, Volume 13 Issue 4

Publisher: ACM Press

Full text available: pdf(2.40 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In a temporal database, time values are associated with data item to indicate their periods of validity. We propose a model for temporal databases within the framework of the classical database theory. Our model is realized as a temporal parameterization of static relations. We do not impose any restrictions upon the schemes of temporal relations. The classical concepts of normal forms and dependencies are easily extended to our model, allowing a suitable design for a database scheme. We pr ...

30 A semiconductor wafer representation database and its use in the PREDITOR



process editor and statistical simulator

D. M. H. Walker, C. S. Kellen, A. J. Strojwas

June 1991 **Proceedings of the 28th conference on ACM/IEEE design automation DAC '91**

Publisher: ACM Press

Full text available: pdf(667.51 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

31 Thoughts on database research: A user perspective



Kurt Ingenthron

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3

Publisher: ACM Press

Full text available: pdf(103.60 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

The future of computer aided design is in object oriented programming. If the database community hopes to participate in this future, it must reexamine some basic assumptions about the architecture of database systems. Database system functionality can be added to object systems but if the performance cost is too high, it will never survive. Below are some suggestions for what can be done at a reasonable performance cost. The object oriented paradigm provides a more practical app ...

32 Federated database systems for managing distributed, heterogeneous, and autonomous databases



Amit P. Sheth, James A. Larson

September 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 3

Publisher: ACM Press

Full text available: pdf(5.02 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A federated database system (FDBS) is a collection of cooperating database systems that are autonomous and possibly heterogeneous. In this paper, we define a reference architecture for distributed database management systems from system and schema viewpoints and show how various FDBS architectures can be developed. We then define a methodology for developing one of the popular architectures of an FDBS. Finally, we discuss critical issues related to developing and operating an FDBS.

33 Effective page refresh policies for Web crawlers



Junghoo Cho, Hector Garcia-Molina

December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4

Publisher: ACM Press

Full text available: pdf(345.52 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this article, we study how we can maintain local copies of remote data sources "fresh," when the source data is updated autonomously and independently. In particular, we study the problem of *Web crawlers* that maintain local copies of remote Web pages for Web search engines. In this context, remote data sources (Websites) do not notify the copies (Web crawlers) of new changes, so we need to periodically *poll* the sources to maintain the copies up-to-date. Since polling the sources ...

Keywords: Web crawlers, page refresh, web search engines, world-wide web

34 Research track: CLOSET+: searching for the best strategies for mining frequent closed itemsets



Jianyong Wang, Jiawei Han, Jian Pei

August 2003 **Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '03**

Publisher: ACM Press

Full text available: pdf(492.93 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Mining frequent closed itemsets provides complete and non-redundant results for frequent pattern analysis. Extensive studies have proposed various strategies for efficient frequent closed itemset mining, such as depth-first search vs. breadthfirst search, vertical formats vs. horizontal formats, tree-structure vs. other data structures, top-down vs. bottom-up traversal, pseudo projection vs. physical projection of conditional database, etc. It is the right time to ask "*what are the pros and c ...*"

Keywords: *association rules, frequent closed itemsets, mining methods and algorithms*

35 Frame-sliced partitioned parallel signature files



Fabio Grandi, Paolo Tiberio, Pavel Zezula

June 1992 **Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '92**

Publisher: ACM Press

Additional Information:

Full text available:  [pdf\(1.14 MB\)](#)

[full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The retrieval capabilities of the signature file access method have become very attractive for many data processing applications dealing with both formatted and unformatted data. However, performance is still a problem, mainly when large files are used and fast response required. In this paper, a high performance signature file organization is proposed, integrating the latest developments both in storage structure and parallel computing architectures. It combines horizontal and vertical app ...

36 [Technology strategy and management: Beware the lure of the horizontal](#)




Michael Cusumano

July 2003 **Communications of the ACM**, Volume 46 Issue 7

Publisher: ACM Press

Full text available:  [pdf\(51.73 KB\)](#)

 [html\(13.75 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

The correct choice of market segmentation can determine product success and help a company rise above the competition.


37 [Declarative programming of graphical interfaces by visual examples](#)



Ken Miyashita, Satoshi Matsuoka, Shin Takahashi, Akinori Yonezawa, Tomihisa Kamada

December 1992 **Proceedings of the 5th annual ACM symposium on User interface software and technology UIST '92**

Publisher: ACM Press

Full text available:  [pdf\(1.26 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Graphical user interfaces (GUI) provide intuitive and easy means for users to communicate with computers. However, construction of GUI software requires complex programming that is far from being intuitive. Because of the "semantic gap" between the textual application program and its graphical interface, the programmer himself must conceptually maintain the correspondence between the textual programming and the graphical image of the resulting interface. Instead, we propose a pr ...

Keywords: constraints, direct manipulation, graphical user interface, layouts, programming by example, visualization

38 [A piggyback method to collect statistics for query optimization in database management systems](#)

Qiang Zhu, Brian Dunkel, Nandit Soparkar, Suyun Chen, Berni Schiefer, Tony Lai

November 1998 **Proceedings of the 1998 conference of the Centre for Advanced Studies on Collaborative research CASCON '98**

Publisher: IBM Press

Full text available:  [pdf\(328.82 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A database management system (DBMS) performs query optimization based on statistical information about data in the underlying data-base. Out-of-date statistics may lead to inefficient query processing in the system. Existing solutions to this problem have some drawbacks such as heavy administrative burden, high system load, and tardy updates. To overcome these drawbacks, our new approach, called the piggyback method, is proposed in this paper. The key idea is to piggyback some additional retriev ...


Keywords: access method, cost estimation, database management system, piggyback analysis, query optimization, statistics collection

39 Complete answer aggregates for treelike databases: a novel approach to combine querying and navigation

Holger Meuss, Klaus U. Schulz

April 2001 **ACM Transactions on Information Systems (TOIS)**, Volume 19 Issue 2

Publisher: ACM Press

Full text available:  pdf(356.60 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The use of markup languages like SGML, HTML or XML for encoding the structure of documents or linguistic data has lead to many databases where entries are adequately described as trees. In this context querying formalisms are interesting that offer the possibility to refer both to textual content and logical structure. We consider models where the structure specified in a query is not only used as a filter, but also for selecting and presenting different parts of the data. If answers are formaliz ...


Keywords: SGML, XML, answer presentation, information retrieval, logic, query languages, semistructured data, structured documents, tree databases, tree matching

40 Correctness of query execution strategies in distributed databases

S. Ceri, G. Pelagatti

December 1983 **ACM Transactions on Database Systems (TODS)**, Volume 8 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.93 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A major requirement of a Distributed DataBase Management System (DDBMS) is to enable users to write queries as though the database were not distributed (distribution transparency). The DDBMS transforms the user's queries into execution strategies, that is, sequences of operations on the various nodes of the network and of transmissions between them. An execution strategy on a distributed database is correct if it returns the same result as if the query were applied to a nondistributed datab ...

Keywords: correctness of database access, distributed database access, read-only transactions, relational algebra

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)